

A Forward-Looking High-Resolution GPR System

Joel Kositsky, * Peyman Milanfar*
SRI International, Menlo Park, CA 94025

ABSTRACT

A high-resolution ground penetrating radar (GPR) system was designed to help define the optimal radar parameters needed for the efficient standoff detection of buried and surface-laid antitank mines. The design requirements call for a forward-looking GPR capable of detecting antitank mines in a 5 to 8 meter wide swath, 7 to 60 meters in front of a mobile platform.

The system has a resolution goal of 15 cm both in range and azimuth. The range and azimuthal resolutions are achieved by using a 2.7 GHz bandwidth and a 4 meter synthetic aperture, respectively. The system uses a fully coherent homodyne stepped-frequency approach with a modulation scheme that produces range dependent power gain to partially offset range losses. Transmit power of 1 to 10 W is available over the entire band, and a large effective dynamic range was built into the receiver. The antennas are mounted on separate transmit and receive computer-controlled high-precision linear drives for creating the synthetic aperture. A data scan entails stepping through all the frequencies, polarizations, and antenna positions before the van is driven forward for the next scan. Preliminary data, the resulting images, and preliminary work on automatic target detection will be presented.

Keywords: GPR, radar, landmine, clutter, mobile, SAR, mine detection, antitank

1. INTRODUCTION

Landmine detection in civilian and military arenas is both a pressing and a challenging task. Standoff detection with high probability of detection (Pd) and low probability of false alarm (Pfa) is especially difficult. In this paper, we outline the design of a very versatile GPR that has been built to help define the radar parameters most amenable to the efficient detection and discrimination of buried and surface-laid antitank mines. We have just begun to take data with the system, and have developed clutter reduction algorithms tested, to date, only on clutter models. Once we have collected and imaged the data from several test (surrogate) minefields, we plan to hone these algorithms to aid in decreasing the Pfa while maintaining a minimal 94% Pd.

The radar engineer faces many tradeoffs in selecting appropriate radar design parameters, such as frequency range, aperture size, system resolution, signal-to-noise ratio (SNR), polarization, and so on. In the first phase of the program, we have elected to implement a very versatile radar system. The data collected in this first phase will be analyzed to help us determine which parameters and analysis algorithms are the most useful in attacking the standoff antitank mine detection problem. In the next phase of this work, we will implement the most efficient subset of the phase 1 GPR, which, together with the most efficacious algorithms, will support the best real-time performance.

The second phase of this work will be the implementation of the most efficient subset of the phase 1 GPR, which, together with the most efficacious algorithms, will support the best real-time performance.

2. BASIC RADAR DESIGN

Our general design guidelines for the phase 1 ground penetrating radar included the need for

- 1) a versatile GPR collection system,
- 2) high resolution image products,
- 3) minimizing radar clutter, and
- 4) a feasible transition to a later phase real-time system.

* Correspondence: E-mail: joel.kositsky@sri.com; Telephone: 650 859-6035; Fax: 650 859-5149; peyman.milanfar@sri.com; Telephone: 650 859-3745; Fax: 650 859 6259

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE APR 1999		2. REPORT TYPE		3. DATES COVERED 00-04-1999 to 00-04-1999	
4. TITLE AND SUBTITLE A Forward-Looking High-Resolution GPR System				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California Santa Cruz,Electrical Engineering Department,1156 High Street,Santa Cruz,CA,95064				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

A fully coherent, continuous wave (CW), stepped-frequency, synthetic aperture radar (SAR) approach was chosen as being the most appropriate and versatile radar implementation. The design is relatively simple and amplifiers are readily available. This approach allows for the production of a flat, wideband spectrum, ease in signal processing and system calibration, and ease in skipping narrowband interferers and reserved frequencies.

Figure 1 shows a block diagram of the basic radar architecture. The stepped CW is synthesized and then modulated (gated) before being amplified, and radiated in the polarization selected by a fast switching network. A similar switching network selects the receive antenna polarization and passes the signal to a homodyne receiver where it is gated, filtered, amplified, mixed down to complex baseband. The baseband signals are passed through an anti-aliasing filter and sampled by two analog-to-digital converters (ADCs) before being read by the system computer.

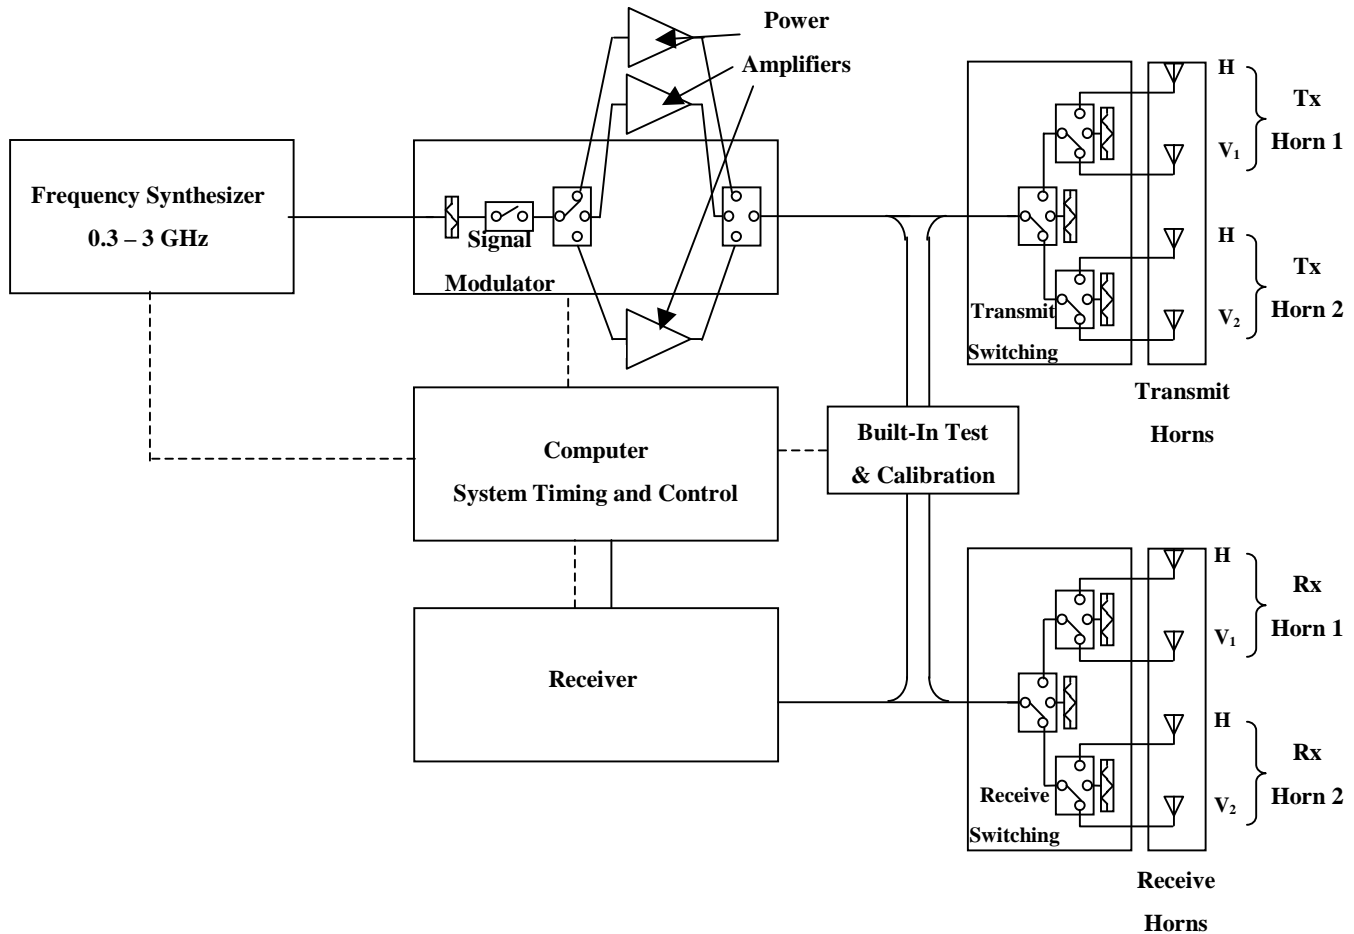


Figure 1 Radar Block Diagram

In this phase of the work, the van will be stationary when data are being collected. A single I/Q sample pair is collected during each pulse repetition interval (PRI) as illustrated in Figure 2. During each PRI, a transmitted wave train at a given frequency is gated to 350 ns; 50 ns later, the receiver is gated open for 350 ns, and after an appropriate delay, the signal is sampled by the digitizer.

The data collection process is illustrated in Figure 3. A programmed number of pulses ("reps") can be transmitted at each polarization and frequency to increase the SNR. A new antenna polarization and/or antenna element is switched in, and the process is repeated. When all combinations are completed, the frequency is incremented to the next step and the entire process is repeated. When all the frequencies have been sampled, the antennas are mechanically stepped to the next positions until data have been collected for the filled synthetic aperture.

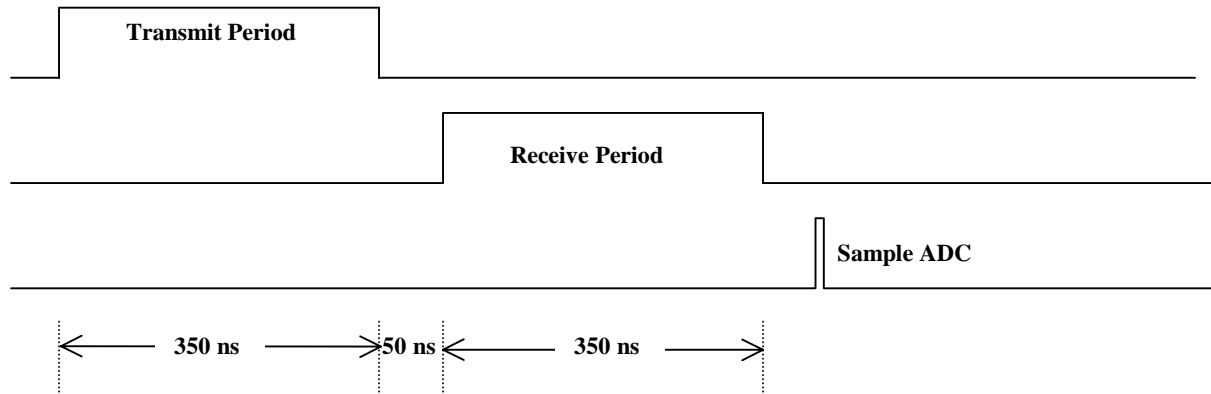


Figure 2 PRI Timing

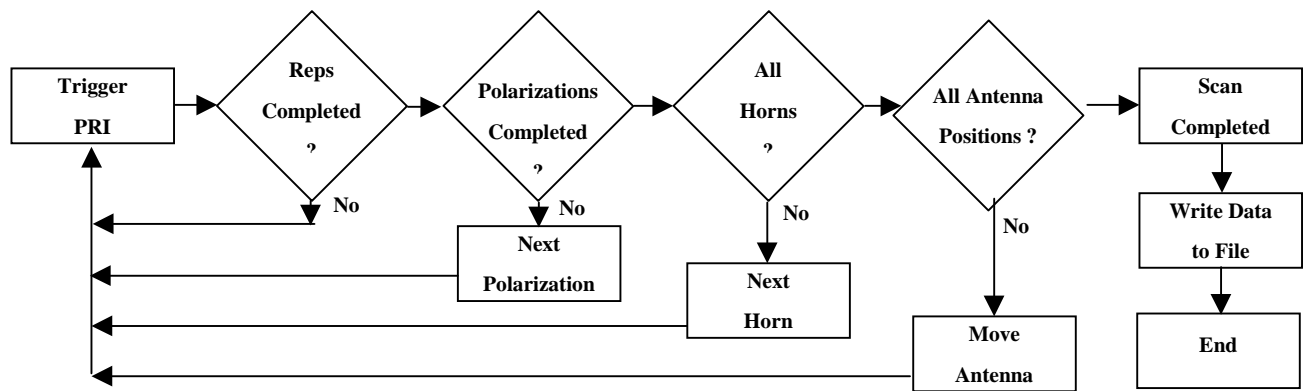


Figure 3 Radar Scan Flowchart

Since this radar is forward-looking, the synthetic aperture is formed by having antenna elements move perpendicular to the direction of travel.[†] Two 1 m, dual-linear polarization, quad-ridged horns are used for each of the transmit and receive antennas. Each pair of antennas is mounted on a computer-controlled linear drive which mechanically moves the horns sideways to produce the synthetic aperture array during data collection. The antennas are mounted one above the other on the top front of the radar van.

Two types of antenna scans are implemented: a standard SAR collection mode, where the transmit and receive horns travel together in the cross-range direction; and a super-SAR collection, where the transmit and receive horns move separately and data are collected at all the combinations and permutations of horn positions at the selected step size.

Data from a single SAR scan are combined in post-processing to form an image of the rectangular area 5 to 8 meters wide, 7 to 60 m in front of the van (Figure 4). The extra data from the super-SAR scan will provide the option of also using super resolution and other imaging techniques that may improve clutter rejection. The radar vehicle is then driven forward (a few meters), stopped, and a new radar scan is performed. The multilook images thus formed can be combined to improve resolution and lower clutter.

[†] The future implementations, as electronically switched array or a real aperture will be used.




Figure 4 Radar Collection Geometry

The basic radar parameters are summarized in Table 1.

Table 1 Radar Parameters

Radar approach:	Stepped frequency SAR
Transmitter waveform:	Chopped CW
Peak power:	10 Watts
Duty cycle:	8%
Frequency span:	0.3 to 3.0 GHz
Frequency step size:	Variable, 0.5 MHz typical
Range:	7 to 60 m
Cross range:	8 m
Azimuthal resolution at 7 m:	15 cm (range dependent)
Range resolution:	15 cm (range independent)
Receiver:	Homodyne/IQ
Digitization:	12 bits (nominal)
Polarization:	HH, VV, HV, VH
Antenna elements:	Quad-ridged horns
Antenna gain:	7 to 12 dBi
Antenna arrays:	1 Tx, 1 Rx (independent)
Array orientation:	4 m horizontal
Antenna scanning:	Mechanical/electrical

The following, more detailed, design description follows the RF path through the system, as shown in Figure 2.

3. DETAILED DESIGN

3.1. Frequency synthesizer

A high quality frequency synthesizer produces the 0.3 to 3.0 GHz RF signals that are modulated, amplified, and radiated by the transmit antennas (see Figure 2). It also produces the local oscillator (LO) signal that is mixed with the received signal for down-conversion in the receiver. The synthesizer is fully programmable, through a parallel interface, with a switching and settling time of 20 μ s or less. Because of the number of frequency steps per data scan, fast synthesizer stepping time is crucial to this project.

3.2. Transmit signal modulator

Simple pulse modulation (switching) of the transmit and receive chains are used to protect the sensitive receiver from saturation during transmission. The planned switch timing is shown in Figure 2, while the radar geometry is shown in Figure 4. The transmit signal is left on for 350 ns, the two-way time-of-flight from the nearest to the farthest range. The receiver is turned on after approximately 50 ns, at the moment when the strong echo from a near-in target at 7 m has almost completely passed the receive antenna, and the weaker echo from a distant target at 60 m is just arriving at the antenna. This timing has the effect of longer exposure time of the receiver to the farther (and therefore weaker) target echoes. The result is a range-dependent signal integration, which helps compensate for the range decrease in reflected power.¹ The receiver is then turned off, while the transmitter is switched on to give another pulse at the same frequency. This is continued as many times as needed to produce the SNR by direct integration.

This modulation scheme has the great advantage of producing inherently stable range-dependent gain, and so increase the effective dynamic range of the receiver. Mathematically, the gain function can be expressed as the convolution between the transmit and receive signals. For rectangular modulation, the convolution gives a sensitivity window that is a linear function of range.

3.3. Power amplifiers and switching

A total of three power amplifiers are used, each operating over about 1 GHz of the frequency range. The lower frequency amplifiers (0.3 to 2 GHz) output 10 W each, while the highest frequency amplifier (2 to 3 GHz) develops 1 W. This diminished power output in the highest frequency band is partially compensated for by the increased directivity of the horns for that range. The computer-controlled electronic switches select the appropriate amplifier and pass the amplified signal to the antenna switching units.

The polarization and antenna element selection-switching unit is a collection of high-power, fast (several ns) electronic switches which control the polarization and antenna element (horn) chosen to radiate the energy. Because there are two horns per antenna, each with vertical and horizontal polarization capabilities, there are a total of four possible transmit choices. A similar switching box is used on the receive side to handle the four possible receive choices. All switches provide a 50 Ω termination in the off condition.

3.4. Antenna elements

A pair of identical quad-ridged horns are used for each antenna (transmit and receive). They are bilinearly polarized with 1 m² apertures. The horns were measured to have a VSWR of better than 3:1 over most of the frequency band, and a -3 dB half-beamwidth that varies from 22° (at 300 MHz) down to about 6° above 2 GHz.

Use of the correct geometrical factors is of paramount importance in producing focused, accurate images from the radar data. Whereas the position of the physical antenna is easily defined, the position from which the EM radiation seems to originate (the phase center) is needed for data processing. For a horn antenna, the position of the phase center varies with frequency; and, to a lesser extent, with azimuthal and elevation angles from the boresight. We determine the phase center experimentally from test range measurements of the horns. These results are then checked against radar data returns produced by a strong point target.

3.5. Radar receiver

The radar receiver is a narrowband, tunable, homodyne receiver covering the 0.3 to 3.0 GHz operating frequency band. At the front end, 300 MHz high-pass filter and a diode limiter protect the sensitive receiver from damage caused by high power levels coupled directly from the transmit antennas. A low-noise preamplifier boosts the signal 18 dB before the electronic mixer that gates the receive signal. This "blanker" is followed by a computer-controlled step attenuator, and an

optional yttrium-iron-garnet programmable bandpass filter to add further protection and band limiting before the second amplifier. The video leak-through from the various switching transients are blanked by carefully adjusting the relative timing of the pulse gating mixer and the front end switches.

A splitter divides the RF signal before it is passed to two matched, wideband mixers that are driven at 90° to produce I and Q signals. Each of these is passed through a low-pass, anti-aliasing filter and a final amplifier before being digitized by the 12-bit ADCs.

The receiver has a measured noise figure value of less than 5 dB, and a fairly constant gain which is matched to the ADC. The noise floor at the system bandwidth (5 MHz) is about -95 dBm, and the second-order spurious-free dynamic range (SFDR) is 55 dB, while the third-order SFDR is 62 dB.

3.6. Calibration and built-in-test

A module was built to provide a calibration and built-in-test (BITE) capability. This assembly allows part of the transmit signal to be coupled to the receiver through a calibrated delay line to simulate the time of flight of the radar signals. A switching network connects the couplers when desired, while maintaining the specified Tx/Rx isolation. Fixed attenuators in the path adjust the coupled signal to match the receiver input level, while a variable attenuator adjusts the signal over the receiver's complete dynamic range. All switches and variable attenuators are under computer control so a complete BIT sequence can be run automatically and efficiently before and after data collection. The bi-directional couplers also permit the receiver to be used to check all the antenna switching circuits and the reflection characteristics from the horns, in addition to sampling the direct transmit signal.

In order to calibrate the I/Q balance, many samples, at a given frequency, are recorded as the phase angle is varied by driving the LO with a second frequency synthesizer offset slightly from the signal frequency.

3.7. Radar platform

The platform for the GPR is an SRI-owned van that had been converted into a mobile laboratory, complete with built-in shock-mounted equipment racks, 120/240 VAC power generators, air conditioning, work surfaces, and equipment cupboards. A wooden deck is securely bolted to the roof, projecting over the metallic edge of the van, providing an excellent nonconducting base for the antenna assemblies.

The antenna platform was fabricated to be easily installed, lightweight, rigid, and electromagnetically transparent. Tracks and computer-controlled linear drives allow the transmit and receive antennas to be moved independently over a 4 m lateral aperture. A video camera is mounted on this platform to optically display and record the collection area, while a GPS system provides time and location logging capabilities.

3.8. Computer system

The computer system chosen to run the data acquisition and test system and to log the data is a VXI Slot0 controller with a 200 MHz Pentium CPU. The board is housed in an air-cooled rack-mounted VXI chassis along with the timing, the parallel I/O, and the ADC boards.

Two precision timing boards are used for the radar's critical timing functions during each PRI. These boards can produce up to 10 separately programmable delay outputs, or up to 5 delay-and-width outputs. Each module has a 32-bit dynamic range, a 39 ps resolution, low jitter, and a 2.5 MHz repetition rate capability. The delay parameters can be programmed on-the-fly at VME-bus speeds and updated coherently.

The 12 bit, 2-channel digitizer boards are VXI message-based instruments with a maximum sampling speed of 10 MS/s and a 512 kSample FIFO. The 96 bit parallel I/O board is a 12 port digital input/output module used for data communication and digital control, configured in 8-bit units as either command or telltale bits. Some of the command signals are converted to more robust differential signals for controlling the more remote switching equipment on the van roof.

The data acquisition program is written in C operating under a Windows/NT operating system. Using a LabView-based approach, we created a graphical user interface (GUI) that presents the operator with a computer screen containing "buttons," "switches," pull-down menus, and fields for entering numerical and other data. The collection parameters that can be input include the polarizations to be included in the scan; the number of sample loops ("reps"); the start, stop, and frequency step size; the amount of attenuation added to the receive chain as a function of frequency; the synthetic aperture size and synthetic array spacing; and the PRI length. Frequencies to be skipped are entered from a separate interface. An A-scope type display

(Figure 5) lets the operator view the digitized radar returns in real time. Data can also be written to disk for later imaging and processing.




Figure 5 A-Scope Display (4 polarizations, 4 horn combinations, 2 frequencies)

3.9. System timing and synchronization

Figure 2 showed the timing of a single collection datum (PRI). Within each PRI, a single software trigger results in a preprogrammed sequence of events to be triggered by the precision timing hardware:

- 1) The correct transmit antenna and polarization are switched on
- 2) The other antennas and polarizations are terminated
- 3) The transmit pulse is gated by the modulation switch
- 4) The correct receive antenna and polarization are switched on
- 5) The receiver blanking is removed
- 6) The ADC samples the receiver output I and Q channels
- 7) The protective blankers are switched on.

Each PRI lasts a total of 5 μ s, at a typical collection rate of 200 kHz.

All the digital radar equipment is kept coherent by using the high-precision internal ovenized crystal oscillator of the RF synthesizer as the master clock to synchronize the other clocks.

4. POST-PROCESSING

The main steps in our post-processing consist of correcting the data for system anomalies, imaging the data, and then analyzing the images to discriminate the mines from the clutter.

4.1. Image processing

After being transformed to the time domain, the data are imaged by direct integration using a back-projection algorithm. Before being imaged, the data are corrected for system transfer function characteristics, RFI, cable delays, and antenna phase center movement.

4.2. Target Detection Algorithm Development

Our experience with both airborne and ground-based GPRs leads us to be confident in our ability to detect shallow buried antitank mines at significant standoff distances. We recognize, however, that detection alone is only part of the solution to the existing problem. The ability to discriminate among the mines and other man-made and natural objects and other sources of clutter is crucial to the utility and acceptance of an antimine technology.

We are in the process of developing and implementing automatic target detection (ATD) approaches and algorithms that show promise in being able to greatly decrease the false alarm rate without severely lowering the probability of detection. These algorithms are very powerful image analysis tools, while being adaptive and efficient enough to work in a real-time operational scenario.

In the course of this work, several approaches for the automatic target detection and discrimination task have been considered. While it is possible, but rather difficult, to explicitly model the clutter statistics for the images produced by this sensor from first principles, it is more convenient to apply a rather general stochastic modeling framework to the problem. To this end, we have implemented three distinct algorithms for capturing the correlation structure in images and will apply these methods to the problem of background decorrelation (or whitening) of the images. As we are only now beginning data collection with our GPR, most of the algorithm work, to date, has been with synthetic radar data combined with clutter based on existing soil models.

4.2.1. Constant false alarm rate (CFAR) statistical normalization

Statistical normalization is often employed in high signal-to-clutter scenarios to detect targets; it can be quite effective² in reducing the number of false alarms in SAR ATD tasks, and is relatively simple to implement. Briefly, standard statistical normalization reassigns each pixel value in the given image by subtracting the local mean around this pixel from it, and dividing the result by the locally estimated standard deviation. That is, the following test is performed in pixel-wise fashion:

$$\text{Target present if } \frac{f(i, j) - \bar{f}}{\hat{\sigma}} \geq d, \quad (1)$$

Where $\bar{f}, \hat{\sigma}$ are, respectively, the estimated mean and variance for the background, and d is a prescribed threshold to achieve a desirable false alarm rate. Typically, these local statistics are computed in a region that is separated from the center pixel by a “guard-region” so as to avoid including possible target pixels in the computation of the background statistics. While very simple in nature, and sometimes limited in utility in practice, it is nonetheless a powerful processing paradigm as it is naturally adaptive since the normalization parameters are estimated for each neighborhood about a pixel of interest. This algorithm is, therefore, able to deal with changing statistics across the image. On the other hand, being a simple first-order model, it is generally less useful for more complex scenarios where the signal-to-clutter ratio (SCR) is lower or the clutter correlation structure is more complex.

Similar to existing techniques, our current implementation of the CFAR normalization algorithm estimates background mean and variance from an annular area around a pixel. In addition, in an improvement over the standard CFAR processor, our estimates of local background mean and variance are robust. That is to say, instead of directly computing these values, we first rank-order (sort) the values in the selected annular region and trim the tails of this list of numbers to the extent determined by a pair of user-defined parameters. This process further ensures that the estimates are free of contamination from target pixels and other irrelevant values such as padded zeros around the edges of the image. This simple change produces improved results over the standard CFAR processor.

4.2.2. Spatial autoregressive (AR) model

A powerful alternative to the simple normalization approach is spatial AR modeling. In this approach, the local correlation structure of the image is captured with a linear predictive (AR) model.³ The coefficients of the model are estimated from the image data directly using a multidimensional formulation of the Yule-Walker equations. The residual image is then produced by subtracting the measured image pixel values from the predicted pixel values. These residuals, once decorrelated, can be used for detection purposes. Mathematically, if we denote the (lexicographically scanned) vector of image values as F , we have

$$F = \Phi F + Z \quad (2)$$

$$Z^{(w)} = \Lambda^{-1/2} (I - \Phi) F, \quad (3)$$

where $\Phi, Z, \Lambda, Z^{(w)}$ are, respectively, the AR coefficient matrix, the residual vector, the residual covariance matrix, and the whitened residual vector. In the absence of other prior information, the whitened residuals can be compared directly to a threshold to declare detections.

To delineate some differences and parallels with the statistical normalization algorithm described earlier, we note that in the AR approach, the model coefficients are estimated for the entire image, or a particular segment thereof (if given a segmentation map). The relative merits of the two approaches are that the AR model is capable of capturing more complex correlation structure as it uses more parameters in modeling the background. In addition, the numerical implementation of this approach is well understood from linear prediction theory. Another important feature of the AR paradigm is that it can be implemented in multiple dimensions; that is, it is capable of whitening data sets that consist of many (possibly correlated) images of the same scene. In our context, this is particularly attractive since our sensor will produce multipolarization data so that to each scene there will correspond at least three distinct images (HH, HV, and VV polarizations). A target of interest may be better delineated in one channel than in the others, but by exploiting all available data, the correlation structure may be better estimated, hence improving the overall detection results.

4.2.3. Multiscale (MS) AR model

The multiscale model is based on the notion that correlation structure may be exploited and modeled at multiple scales of resolution. The model associates a (not necessarily spatially local) neighborhood of pixels at one scale to a group of pixels at the next coarser scale via a linear transformation. In this sense, the MS model is also autoregressive—in scale, rather than in space. Figure 6 displays examples of two different types of multiscale models. Here we employ the type of model shown on the right.⁴




Figure 6 Multiscale Models

Given an image, the parameters of the multiscale model are estimated by decomposing the image onto the scale "pyramid" and solving a sequence of least-squares problems recursively in scale starting from the finest scale.⁴ These estimated parameters are then used to predict the values of the pixels at a given scale from the pixels at the next coarser scale. The difference between the predicted and actual values then constitutes the residual or detail image. Finally, the residuals at all available scales are normalized according to their respective scale-varying statistics and combined to produce the overall residual (or whitened) image, which can subsequently be compared to a threshold.

4.2.4. Examples with simulated data

After some preliminary analysis, we determined that the AR and MS algorithms are the most promising and suitable for the present application, as they appear significantly more effective and reliable in modeling the stochastic nature of the background noise/clutter. Therefore, in this section we briefly present some preliminary result using only these techniques on simulated data.

In Figure 7, we display a simulated SAR image of 18 point targets: on the left the uncluttered “clean” SAR image is displayed, and on the right the cluttered image at a peak SNR of 3 dB is displayed. The cluttered image was obtained by adding spatially correlated Gaussian noise to the complex pixel values and taking the magnitude of the result. The spatial correlation structure of the noise is simulated using a Gaussian kernel with a 7×7 support. The cluttered image in Figure 8 was then used to test the AR and MS algorithms, as described earlier. After whitening, the images are thresholded at the 99% confidence interval and the resulting (binary) detection images are morphologically processed to remove isolated points and to cluster nearby detections, finally producing the images in Figure 8. As can be observed, at least qualitatively, most of targets have been detected successfully with few false-alarm pixels. While this is only a preliminary result on simulated data, we are encouraged that it will perform well on real data. Much work remains to be done to prove this assertion.




Figure 7 Images from Synthetic Data




Figure 8 Clutter Reduction Using Two Algorithms

ACKNOWLEDGMENTS

This research project is funded under contract number DAAB07-98-C-6007 with the CECOM, Night Vision Electronic Sensor Division, Mine Neutralization Branch, Fort Belvoir.

REFERENCES

1. S.-E. Hamran, D. T. Gjessing, J. Hjelmstad, and E. Aarholt, "Ground penetrating synthetic pulse radar: dynamic range and modes of operation," *J. Applied Geophysics*, **33**, pp. 7-14, 1995.
2. L. Novak, G. J. Owirka, and C. M. Netishen, "Performance of a high-resolution polarimetric SAR automatic target recognition system," *The Lincoln Laboratory Journal*, **6**, 1, pp. 11-24, 1993.
3. D. M. Silva, I. E. Abdou, and R. E. Warren, "Optimum detection of small targets in a cluttered background," *Optical Engineering*, **37**, 1, pp. 83-92, January 1998.
4. P. Milanfar, R. R. Tenney, and A. S. Willsky, "Modeling and estimation for a class of multiresolution random fields," *Proceedings of the International Conference on Image Processing*, pp. 397-401, 1994.